arXiv:1512.05698v1 [stat.ML] 17 Dec 2015

Oracle inequalities for ranking and U-processes with Lasso penalty

Wojciech Rejchel
Faculty of Mathematics and Computer Science
Nicolaus Copernicus University
Chopina 12/18, 87-100 Toruń, Poland
Faculty of Mathematics, Informatics and Mechanics
University of Warsaw
Banacha 2, 02-097 Warsaw, Poland
wrejchel@gmail.com

Abstract

We investigate properties of estimators obtained by minimization of U-processes with the Lasso penalty in high-dimensional settings. Our attention is focused on the ranking problem that is popular in machine learning. It is related to guessing the ordering between objects on the basis of their observed predictors. We prove the oracle inequality for the excess risk of the considered estimator as well as the bound for the l_1 distance $|\hat{\theta} - \theta^*|_1$ between the estimator and the oracle.

Keywords: high-dimensional problem, machine learning, oracle inequality, penalized risk minimization, sparse model, U-process.

1 Introduction

Model selection is an important challenge, if one works with data sets containing many predictors. Finding relevant variables helps to understand better the problem and improves statistical inference. In the literature there are many methods solving such problems. One of them is empirical risk minimization with the penalty, for instance Lasso (Tibshirani, 1996). The main characteristic of this procedure is an ability to select relevant predictors and estimate unknown parameters simultaneously. In the paper we apply these ideas to the pairwise ranking problem (ordinal regression) that

relates to predicting or guessing the ordering between objects on the basis of their observed predictors. The problem of ranking has numerous applications in practice, for instance in information retrieval, banking or quality control. This is one of the reasons why it has been extensively studied recently (Freund et al., 2004; Cortes and Mohri, 2004; Agarwal et al., 2005; Cossock and Zhang, 2006; Rudin, 2006; Clémençon et al., 2008; Rejchel, 2012; Chen and Wu, 2013; Lai et al., 2013; Laporte et al., 2014). The ranking problem is one of the examples of using U-processes in machine learning. The goal of the paper is to investigate properties of estimators obtained by minimization of U-processes with the Lasso penalty. We study the quality of such estimators in prediction as well as model selection. We focus our attention on ranking estimators, but our results can be easily generalized as we will explain below. We start the paper with describing the statistical framework of the ranking problem.

Let us consider two objects that are randomly selected from a population. We assume that they are described by a pair of independent and identically distributed (with respect to the measure P) random vectors $Z = (X, Y)$ and $Z' = (X', Y')$ taking values in $\mathcal{X} \times \mathcal{Y}$, where $\mathcal{X} \subset \mathbb{R}^d$ and $\mathcal{Y} \subset \mathbb{R}$. Random vectors X and X' are regarded as predictors, while Y and Y' are unknown variables that define the ordering between objects. Namely, the object z is "better" (faster, stronger etc.) than the object z', if $y > y'$. Our task is to predict the ordering between objects only on the basis of observations X, X'. To do it we construct functions $f : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$, called ranking rules, that predict the ordering between objects in the following way:

$$\text{if } \ f(x, x') > 0, \text{ then we predict that } \ y > y'.$$

The natural approach is to look for a function f which, for a fixed loss function $\phi : \mathbb{R} \to [0, \infty)$, minimizes the theoretical risk

$$Q(f) = \mathbf{E}\phi\left[\text{sign}(Y - Y')f(X, X')\right] \tag{1}$$

in some family of ranking rules $\mathbf{F}$ and $\text{sign}(t) = 1$ for $t > 0$, $\text{sign}(t) = -1$ for $t < 0$ and $\text{sign}(t) = 0$ for $t = 0$. We cannot calculate $f^0 = \arg\min_{f \in \mathbf{F}} Q(f)$ directly, but if we possess a learning sample of independent and identically distributed (with respect to the measure P) random vectors $Z_1 = (X_1, Y_1), \ldots, Z_n = (X_n, Y_n)$, then we can minimize the empirical analog of the risk (1). In the ranking problem a natural candidate for the empirical risk is a U-statistic

$$Q_n(f) = \frac{1}{n(n-1)} \sum_{i \neq j} \phi_f(Z_i, Z_j), \tag{2}$$

where $\phi_f(z, z') = \phi[\text{sign}(y - y')\, f(x, x')]$. The sum in (2) is taken over all pairs of distinct indices $(i, j) \in \{1, \dots, n\}^2$. Obviously, we could also divide a sample into two halves and consider the unbiased estimator of the risk (1) as an average of independent random variables

$$\frac{1}{N} \sum_{i=1}^{N} \phi_f \left(Z_i, Z_{N+i} \right), \tag{3}$$

where $N = \left\lfloor \frac{n}{2} \right\rfloor$. We could also divide a sample into even and odd observations and so on. However, the U-statistic (2) is an average over all such divisions, that is for every permutation π of a set $\{1, \dots, n\}$ we have

$$\frac{1}{n(n-1)} \sum_{i \neq j} \phi_f(Z_i, Z_j) = \frac{1}{n!} \sum_{\pi} \frac{1}{N} \sum_{i=1}^{N} \phi_f \left(Z_{\pi(i)}, Z_{\pi(N+i)} \right).$$

Thus, the esimator (2) seems to be more "stable" than (3). Indeed, provided that mild conditions are satisfied the estimator (2) has least variance among unbiased estimators of (1), see Serfling (1980).

The family $\mathcal{F}$, that the estimator of f^0 will come from, is a class of linear combinations of a given collection of functions. Namely, consider a finite (but usually very large) family of base functions $\{\psi_1, \dots, \psi_m\}$, where $\psi_k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}, k = 1, \dots m$. We focus on the family

$$\mathcal{F} = \left\{ f_\theta(x, x') = \sum_{k=1}^{m} \theta_k \psi_k(x, x') : \theta \in \Theta \right\} \subset \mathbf{F}, \tag{4}$$

and the set Θ is a convex subset of $\mathbb{R}^m$. If $m = d$ and $\psi_k(x, x') = x_k - x'_k$, then $\mathcal{F}$ is a class of linear ranking rules. In this case we can think that the number of predictors d that we observe is huge, but only a few of them are significant in the model. In this sense we say that the model is "sparse".

Furthermore, in high-dimensional problems one usually adds the penalty to the empirical risk (2). In the paper we consider the Lasso penalty, so we are to minimize, over the family (4), a function

$$Q_n(f_\theta) + \hat{\lambda}_n |\theta|_1, \tag{5}$$

where $|\cdot|_1$ is the l_1-norm of a vector θ, i.e. $|\theta|_1 = \sum\limits_{k=1}^{m} |\theta_k|$, and $\hat{\lambda}_n > 0$ is a smoothing parameter that will be estimated from the data, as it will be explained later. This parameter is a balance between minimizing the

empirical risk (2) and the penalty. We denote the minimizer of (5) by $f_{\hat{\theta}}$. The form of the penalty is crucial, because its singularity at the origin implies that some coordinates of the minimizer are exactly zeros, if $\hat{\lambda}_n$ is sufficiently large. Thus, by minimizing the function (5) we simultaneously select significant parameters in the model and estimate them. Finally, we assume that the model is "sparse", that is using only few base functions we can appropriately approximate the best function f^0 with respect to the risk.

The $0-1$ loss function $\phi(t) = \mathbb{I}_{(-\infty,0)}(t)$ seems to be the most adequate, and the theoretical risk (1) becomes probability of incorrect ranking. However, in this case the empirical risk (2) is discontinuous and minimization of the function (5) is computationally difficult. Thus, we assume that the loss fuction ϕ is convex which makes the function (5) convex with respect to θ.

In the literature one can find many papers concerning properties of Lasso estimators. They investigate the risk of estimators as well as model selection ability (Greenshtein, 2006; van de Geer and Tarigan, 2006; Zhao and Yu, 2006; Meinshausen and Buhlmann, 2006; Bunea et al., 2007; van de Geer, 2008; Bühlmann and van de Geer, 2011; Negahban et al., 2012; Huang and Zhang, 2012). However, these papers study the case that one minimizes an empirical process (the sum of independent random variables) with the Lasso penalty. In the current paper we are to minimize (5), that is a U-process with the Lasso penalty, and we want to study properties of estimators based on such minimization. Obviously, U-processes are more complex than empirical processes, in fact they are generalization of empirical processes. Characteristics of ranking estimators with the penalty were investigated, among others, in Clémençon et al. (2008); Rejchel (2012); Chen and Wu (2013), but these results consider only their quality in predicting the ordering between objects. In many practical problems (in genetics or biology) finding a (small) subset of significant predictors is equally as important (or even more) as predicting accurately the ordering between objects. Therefore, in the current paper we study both prediction and model selection quality of estimators obtained by minimization of (5). The practical performance of such estimators was studied, among others, in Lai et al. (2013); Laporte et al. (2014) on the basis of popular data sets. Thus, we focus on theoretical properties of ranking estimators and prove similar theorems to those contained in papers that are cited at the beginning of this paragraph.

In our main result (Theorem 1) we consider the excess risk of the estimator, that is $Q(f_{\hat{\theta}}) - Q(f^0)$. It is a difference between the risk of the estimator and the risk of the best ranking rule in the class $\mathbf{F}$. We show that the estimator $f_{\hat{\theta}}$ behaves almost like the "oracle" that knows which parameters (base functions) should be included in the model by balancing the approximation

risk and sparseness. Moreover, in the same theorem we show the inequality for the l_1-distance $|\hat{\theta}-\theta^*|_1$ between the estimator and the oracle that concerns model selection quality of the procedure. Our results are nonasymptotic and allow the number of parameters m to grow polynomially fast with the sample size n. Theorem 1 is a strict analog of van de Geer and Tarigan (2006, Theorem 2.2) or Bühlmann and van de Geer (2011, Theorem 6.4), but it relates to U-processes with the Lasso penalty instead of empirical processes.

We have already mentioned that the ranking problem with the Lasso penalty is contained in the general problem of minimizing a U-process with the Lasso penalty. U-processes are used in many problems of statistical learning, for instance in generalized regression models or survival analysis (Han, 1987; Sherman, 1993; Klein and Spady, 1993; de la Peña and Giné, 1999; Bobrowski and Łukaszuk, 2012). Properties of ranking estimators with the Lasso penalty that are described in Theorem 1 can be generalized to U-processes with the Lasso penalty. Indeed, for a function $f : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ we define its risk as

$$\tilde{Q}(f) = \mathbf{E}\, \tilde{\phi}\left[f(X, X'), Y, Y'\right],$$

where $\tilde{\phi} : \mathbb{R} \times \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ is the loss function and $Z = (X, Y), Z' = (X', Y') \in \mathcal{X} \times \mathcal{Y}$ are independent random vectors with the same distribution P. Moreover, assume that random vectors $Z_1 = (X_1, Y_1), \ldots, Z_n = (X_n, Y_n)$ are independent copies of them. Define families $\mathbf{F}, \mathcal{F}$ as above and introduce a notation $\tilde{\phi}_f(z, z') = \tilde{\phi}[f(x, x'), y, y']$. Then the unknown function $f^0 = \arg\min_{f \in \mathbf{F}} \tilde{Q}(f)$ can be estimated by the minimizer $f_{\hat{\theta}}$ of a U-process with the Lasso penalty

$$\tilde{Q}_n(f_\theta) + \hat{\lambda}_n |\theta|_1, \tag{6}$$

where

$$\tilde{Q}_n(f_\theta) = \frac{1}{n(n-1)} \sum_{i \neq j} \tilde{\phi}_{f_\theta}(Z_i, Z_j).$$

Results stated for ranking estimators in the rest of the paper can be easily extended to minimizers of (6).

The paper is organized as follows: in Section 2 we state the main theorem of the paper and discuss it. Its assumptions are elaborated in Section 3. The proof of Theorem 1 is divided into two parts that are contained in Appendices A and B. We conclude the paper in Section 4.

2 Assumptions and main result

We start this section with presenting and briefly discussing conditions that are needed to obtain the results. We explain them more precisely in the next section.

Assumption 1. The loss function ϕ is convex, besides it satisfies the Lipschitz condition with the constant $L > 0$, that is

$$|\phi_{f_\theta}(z, z') - \phi_{f_{\theta'}}(z, z')| \leq L\,|f_\theta(x, x') - f_{\theta'}(x, x')| \tag{7}$$

for every $\theta, \theta' \in \Theta$ and $z = (x, y), z' = (x', y') \in \mathcal{X} \times \mathcal{Y}$. We have already mentioned that convexity of the loss function is needed to make the procedure computationally effective. However, it is also substantially used in the proof of Theorem 1, for instance "to localize the problem" (we can investigate only the neighbourhood of the oracle).

Assumption 2. Let $\mathbf{F}$ be a family of functions with a pseudonorm $||\cdot||$. There exists a strictly convex function $G : [0, \infty) \to [0, \infty)$ with $G(0) = 0$, such that for every $f_\theta \in \mathcal{F}$ we have

$$Q(f_\theta) - Q(f^0) \geq G\left(||f_\theta - f^0||\right). \tag{8}$$

Similar conditions that uniformly bound the excess risk are often called "margin conditions" in the literature (Mammen and Tsybakov, 1999; van de Geer, 2008; Bühlmann and van de Geer, 2011). They are necessary to obtain "fast rates" for estimators based on empirical processes (Mendelson, 2002; Bartlett et al., 2005, 2006) as well as U-processes (Clémençon et al., 2008; Rejchel, 2012). Fast rates refer to probabilistic inqualities that bound the estimation risk by the expression depending on the sample size n like $n^{-\beta}$ with $\beta > \frac{1}{2}$. Furthermore, notice that Assumption 2 (as well as Assumption 3, below) depends on the choice of the pseudonorm $||\cdot||$.

For the function G given in Assumption 2 we define its convex conjugate as a function $H : [0, \infty) \to [0, \infty)$ given as

$$H(v) = \sup_{u \geq 0}\{uv - G(u)\}$$

for every $v \geq 0$. For instance, if $G(u) = au^{\frac{2}{\alpha}}$ for a positive constant a and $\alpha \in (0, 1]$, then $H(v) = \frac{2-\alpha}{2}\left(\frac{\alpha}{2a}\right)^{\frac{\alpha}{2-\alpha}} v^{\frac{2}{2-\alpha}}$. The function H will be involved in a bound for the estimation risk of the estimator and will determine the rate of the oracle inequality.

Before we state the next assumption we need a few notations and a definition. Let S be a subset of $\{1, \ldots, m\}$. We denote by $|S|$ the number of

elements of the set S and its complement by $S' = \{1, \ldots, m\} \backslash S$. Moreover, let θ_S be the vector that equals θ on the set S and has zeros elsewhere, i.e. $(\theta_S)_k = \theta_k \, \mathbb{I}(k \in S), k = 1, \ldots, m$. Thus, we have $\theta = \theta_S + \theta_{S'}$.

Definition 1. *The compatibility condition is fulfilled for the set S with a constant $A(S)$, if for every $\theta \in \mathbb{R}^m$ satisfying $|\theta_{S'}|_1 \leq 3|\theta_S|_1$ the following inequality holds*

$$|\theta_S|_1 \leq \frac{||f_\theta|| \sqrt{|S|}}{A(S)} \, . \tag{9}$$

The above definition is borrowed from Bühlmann and van de Geer (2011, Section 6.7). It constitutes the relation between the pseudonorm $||\cdot||$ and the norm $|\cdot|_1$. The compatibility condition or similar assumptions such as the coherence condition or the restricted eigenvalue condition are standard in high-dimensional problems (Bunea et al., 2007; van de Geer, 2008; Bickel et al., 2009; Bühlmann and van de Geer, 2011).

Assumption 3. The set Θ_1 that minimization in (10), given below, is taken over consists of vectors θ such that the compatibility condition is satisfied for $S_\theta = \{1 \leq k \leq m : \theta_k \neq 0\}$. We will call S_θ the support of the vector θ.

The last technical assumption helps us to handle the stochastic part of the proof of the main result.

Assumption 4. Suppose that

$$\max_{1 \leq k \leq m} |\psi_k|_\infty \leq \sqrt{\frac{n}{\log m}} \, ,$$

where $|g|_\infty = \sup_{x,x'} |g(x, x')|$.

Moreover, we assume that the value

$$C^2 = \max_{1 \leq k \leq m} \mathbf{E} \psi_k^2(X, X')$$

is finite. We do not need to know the value of C^2, the estimator $\hat{C}^2$ of it will be sufficient in further argumentation. We set $\hat{C}^2$ to be the empirical version of C^2, namely it is defined by the following formula

$$\hat{C}^2 = \max_{1 \leq k \leq m} \frac{1}{n(n-1)} \sum_{i \neq j} \psi_k^2(X_i, X_j).$$

Having the estimator $\hat{C}^2$ we can take the smoothing parameter $\hat{\lambda}_n$ in (5) as

$$\hat{\lambda}_n = BL \sqrt{\frac{\log m}{n}} \, \max\left(\hat{C}, 6\right),$$

where B is a universal constant. Its theoretical analog is denoted by

$$\lambda_n = BL\sqrt{\frac{\log m}{n}}\ \max\left(C, 6\right).$$

The form of λ_n and $\hat{\lambda}_n$ are determined by the probabilistic inequality in Theorem 2. We will discuss it later.

Except the above assumptions we will need a few more notations to state the main theorem. The first one is the oracle θ^* that is defined as

$$\theta^* = \arg\min_{\theta\in\Theta_1}\left\{(1+4\delta)\left[Q(f_\theta) - Q(f^0)\right] + 8\delta H\left(\frac{\lambda_n\sqrt{|S_\theta|}}{\delta A(S_\theta)}\right)\right\}, \qquad (10)$$

where δ is an arbitrary number in $\left(0, \frac{1}{4}\right)$ and the set Θ_1 consists of such $\theta \in \Theta$ that the compatibility condition holds for S_θ (see Assumption 3). The oracle is the element from the subclass of the family $\mathcal{F}$ that is able, in the best possible way, to predict the ordering between objects and to select the sparse model simultaneously. Furthermore, set $S_* = S_{\theta^*}, A_* = A(S_*)$ and

$$\varepsilon^* = (1+4\delta)\left[Q(f_{\theta^*}) - Q(f^0)\right] + 8\delta H\left(\frac{\lambda_n\sqrt{|S_*|}}{\delta A_*}\right). \qquad (11)$$

Now we can state the main result of the paper.

Theorem 1. *Suppose that $\delta \in \left(0, \frac{1}{4}\right)$ is an arbitrary number and Assumptions 1-4 are satisfied. Probability that at least one of the inequalities*

$$(1-4\delta)\left[Q(f_{\hat{\theta}}) - Q(f^0)\right] + \hat{\lambda}_n|\hat{\theta} - \theta^*|_1 \le 2\varepsilon^* \qquad (12)$$

or

$$Q(f_{\hat{\theta}}) - Q(f^0) + \hat{\lambda}_n|\hat{\theta} - \theta^*|_1 \le 4\varepsilon^* \qquad (13)$$

is satisfied is at least $1 - \frac{3}{m^2}$.

Theorem 1 states that the excess risk of the estimator can be upper bounded, with respect to the constants, by (11) which is the sum of two expressions. The first one $Q(f_{\theta^*}) - Q(f^0)$ is the approximation risk (the excess risk of the oracle) which describes how well the best rule f^0 can be approximated (with respect to the risk) by the sparse oracle f_{θ^*}. The second term in (11) is the bound for the estimation risk. It depends on the function H from Assumption 2, the value A_* from the compatibility condition and the number of nonzero coordinates of the oracle. In the next section we give examples of interesting loss functions that satisfy Assumption 2 with

$G(u) = au^{\frac{2}{\alpha}}$ for a positive constant a and $\alpha \in (0, 1]$. Notice that in these cases we would obtain the bound for the estimation risk that behaves like

$$\left(\frac{\log m \ |S_*|}{n}\right)^{\frac{1}{2-\alpha}} , \tag{14}$$

if we forget about dependence on values A_*, C^2 and other constants. Therefore, for sparse models we get oracle inequalities with fast rates, if m does not grow exponentially fast with n, that is a standard requirement for high-dimensional problems.

Furthermore, analysing bounds for the excess risk of the estimator given in inequalities (12) or (13) we claim that the estimator behaves almost like the oracle that knows beforehand which parameters (base functions) should be chosen to approximate appropriately the best function f^0. Indeed, suppose that we knew a priori the support S_*. Then we would estimate only parameters contained in the set S_*, while setting others as zeros. The excess risk of such estimator can be bounded by the sum of the approximation risk and the estimation risk (Rejchel, 2012). Omitting constants, the latter is given by $H\left(\sqrt{\frac{|S_*|}{n}}\right)$, that is simply $\left(\frac{|S_*|}{n}\right)^{\frac{1}{2-\alpha}}$ for $G(u) = au^{\frac{2}{\alpha}}$ in Assumption 2. Therefore, the term $\log m$, that appears in (11) or (14), seems to be a price that we have to pay for not knowing a priori which parameters are in the model.

Besides, Theorem 1 gives the probabilistic inequality for the l_1 distance $|\hat{\theta} - \theta^*|_1$ between the estimator $\hat{\theta}$ and the oracle θ^*. This distance is small, if the best function f^0 can be well approximated by the sparse oracle with respect to the risk, i.e. the excess risk $Q(f_{\theta^*}) - Q(f^0)$ is small. In this case the procedure estimates accurately parameters that are significant in the model (contained in the support S_*). But it is rather not able to discard all irrelevant coefficients. The remedy for that could be using the thresholded Lasso (Zhou, 2009) and the obtained bound for the distance $|\hat{\theta} - \theta^*|_1$ would be useful in proving that supports of this modified estimator and the oracle coincide.

The formula of λ_n is the consequence of Theorem 2 that is a probabilistic inequality for a U-process and can be found in Appendix A. Besides, there is a relation between λ_n and the confidence level in Theorem 1 that can be easily seen in proofs in Appendices. Roughly speaking, the confidence level can be increased, but it requires the growth of λ_n, that makes the bound for the estimation risk worse. For simplicity, we choose the confidence level to be $1 - \frac{3}{m^2}$. Besides, the universal constant B, that appears in formulas λ_n or

$\hat{\lambda}_n$, should be greater than 998. Certainly this value can be decreased, but we have not attempted to optimize it nor other constants in the paper. We have focused on rates of inequalities. Finally, it is worth to notice that the smoothing parameter $\hat{\lambda}_n$ can be calculated having the sample. In particular, it does not depend on unknown values C^2 nor A_*. However, in practice the choice of $\hat{\lambda}_n$ based on cross validation is recommended.

Furthermore, Theorem 1 is similar to theorems concerning properties of estimators obtained by minimization of an empirical process with the Lasso penalty, especially those in van de Geer and Tarigan (2006, Theorem 2.2) and Bühlmann and van de Geer (2011, Theorem 6.4). Notice that the empirical process relates to the sum of independent random variables. Our problem is based on the U-process that relates to the sum of dependent random variables. Nevertheless, the results that we obtain in this more intricate scenario strictly correspond to those obtained for empirical processes.

The proof of Theorem 1 is divided into two parts that are contained in Appendices. The first one can be called "stochastic". We state a probabilistic inequality that bounds U-processes in an appropriate way. It is based on quite simple but very helpful concentration inequality for U-processes. The second part of the proof is "deterministic". We adapt standard argumentation concerning empirical risk minimization with the Lasso penalty (van de Geer and Tarigan, 2006; van de Geer, 2008; Bühlmann and van de Geer, 2011) to the ranking problem (U-processes).

3 Discussion on assumptions

The first and the last assumption of Theorem 1 are standard. Thus, we focus on Assumption 2 and the compatibility condition that is involved in Assumption 3. Notice that they both depend on the choice of the pseudonorm $||\cdot||$. We study two cases: the $\mathbb{L}^2$-pseudonorm $||f||_2 = \sqrt{\mathbf{E} f^2(X, X')}$ and the "conditional" pseudonorm $||f||_c = \sqrt{\mathbf{E}\left[\mathbf{E}_{X'} f(X, X')\right]^2}$, where $\mathbf{E}_{X'} f(x, X') = \mathbf{E}\left[f(X, X')|X = x\right]$ is the conditional expectation of the function f. Obviously, for every function f we have $||f||_c \leq ||f||_2$. Therefore, if the condition (8) is satisfied for $||\cdot||_2$, then it is also fulfilled with $||\cdot||_c$ provided that the function G is nondecreasing. On the other hand, the condition (9) with $||\cdot||_c$ implies the same one with $||\cdot||_2$.

3.1 On compatibility condition

Consider the pseudonorm $||\cdot||_2$. Let $\Psi(x,x')$ stands for the vector of base functions $[\psi_1(x,x'),\ldots,\psi_m(x,x')]^T$ for $x,x' \in \mathcal{X}$. Besides, let a matrix $\Sigma = \mathbf{E}\Psi(X,X')\left[\Psi(X,X')\right]^T$ describe dependence between base functions. If the smallest eigenvalue ρ of the matrix Σ is positive, then it is well-known that the condition (9) is satisfied for every set $S \subset \{1,\ldots,m\}$ with the constant $\sqrt{\rho}$. Thus, the compatibility condition, that looks similarly to the smallest eigenvalue requirement, is significantly weaker than that, because the compatibility condition relates only to vectors satisfying $|\theta_{S'}|_1 \leq 3|\theta_S|_1$.

Furthermore, notice that in the linear case, i.e. $f_\theta(x,x') = \theta^T(x-x')$, the matrix Σ is just twice the variance-covariance matrix $\mathbf{Var}(X)$ of the predictor X. Moreover, in this example the difference between $||f_\theta||_2$ and $||f_\theta||_c$ is rather negligible, because $||f_\theta||_2^2 = 2\,\theta^T\mathbf{Var}(X)\theta$ and $||f_\theta||_c^2 = \theta^T\mathbf{Var}(X)\theta$.

3.2 On Assumption 2

We have already mentioned that the convex conjugate H of the function G from Assumption 2 determines the rate of the bound of the estimation risk. We will show that distinct choices of the pseudonorm $||\cdot||$ can lead to significantly different rates. Similar argumentation, but related to the $0-1$ loss function, can be found in (Clémençon et al., 2008, Chapter 5).

Consider the hinge loss function $\phi(t) = \max(0, 1-t)$ and f^0 being the minimizer of the theoretical risk (2) among all measurable ranking rules $f : \mathcal{X}\times\mathcal{X} \to \mathbb{R}$. It is not difficult to calculate that

$$f^0(x,x') = \begin{cases} 1, & \text{if } \eta(x,x') > \frac{1}{2}, \\ -1, & \text{otherwise,} \end{cases}$$

where $\eta(x,x') = \mathbb{P}\left(Y > Y' | X = x, X' = x'\right)$. To simplify calculations we restrict to the standard linear model, that is $Y = \theta_0^T X + \varepsilon$ and $Y' = \theta_0^T X' + \varepsilon'$, where $X, X' \sim N(0,V), \varepsilon, \varepsilon' \sim N\left(0,\sigma^2\right)$ and we consider a subclass $\mathcal{F}_1 = \{f_\theta \in \mathcal{F} : |f_\theta|_\infty \leq 1\}$. Therefore, we can calculate that $\eta(x,x') = \Phi\left(\theta_0^T(x-x')\right)$ and Φ is the distribution function of the standard normal variable. It implies that $f^0(x,x') = \mathrm{sign}\left(\theta_0^T(x-x')\right)$ and

$$Q(f_\theta) - Q(f^0) = \mathbf{E}\left|f_\theta(X,X') - f^0(X,X')\right|\left|2\eta(X,X') - 1\right|.$$

In the considered model it is not difficult to prove that for every $\alpha \in (0,1)$ there exists a positive $B(\alpha)$ such that for every $x \in \mathcal{X}$ we have

$$\mathbf{E}_{X'}\left|2\eta(x,X') - 1\right|^{-\alpha} \leq B(\alpha). \tag{15}$$

We will show that the condition (15) implies, for every $f_\theta \in \mathcal{F}_1$, the inequality

$$Q(f_\theta) - Q(f^0) \geq \left[B(\alpha)2^{2-\alpha}\right]^{-\frac{1}{\alpha}} ||f_\theta - f^0||_c^{2/\alpha} . \tag{16}$$

Indeed, notice that the squared pseudonorm $||f_\theta - f^0||_c^2$ equals

$$\mathbf{E}_X \left[\mathbf{E}_{X'} \left(f_\theta(X, X') - f^0(X, X')\right) \left|2\eta(X, X') - 1\right|^{\alpha/2} \left|2\eta(X, X') - 1\right|^{-\alpha/2}\right]^2 ,$$

that is, by the Cauchy-Schwarz inequality and (15), upper bounded by

$$B(\alpha)\, \mathbf{E} \left|f_\theta(X, X') - f^0(X, X')\right|^2 \left|2\eta(X, X') - 1\right|^\alpha . \tag{17}$$

Moreover, we know that

$$|f_\theta(x, x') - f^0(x, x')|^2 \leq 2^{2-\alpha} |f_\theta(x, x') - f^0(x, x')|^\alpha$$

for every x, x'. Therefore, Jensen's inequality implies that the expression (17) is not greater than

$$B(\alpha)\, 2^{2-\alpha} \left[\mathbf{E} \left|f_\theta(X, X') - f^0(X, X')\right| \left|2\eta(X, X') - 1\right|\right]^\alpha ,$$

which implies the inequality (16). Therefore, in the considered model we can take the function G in Assumption 2 arbitrarily close to the quadratic one. It implies that the function H also can be arbitrarily close to the quadratic one. Thus, the bound of the estimation risk seems to have almost optimal rate. The situation is different, if one considers the pseudonorm $||\cdot||_2$ instead of $||\cdot||_c$. Using above assumptions on the model we can prove that the condition (8) holds with the function $G(u)$ proportional to u^4, so the function H is propotional to $u^{4/3}$. Indeed, it is enough to use van de Geer and Tarigan (2006, Lemma 3.1) and notice that their condition (20) holds in our model with the exponent one. Obviously, the rate of the estimation risk is worse than in the previous case. To the best of our knowledge, it is rather not possible to bring this result much closer to the one with the pseudonorm $||\cdot||_c$.

In the previous paragraph we describe a model that the stipulation (8) is implied by the condition (15) concerning the probabilistic property of the model. However, in many practical examples convexity of the loss function ϕ (in the appropriate sense) is enough to guarantee the assumption (8). Indeed, consider the model with the pseudonorm $||\cdot||_2$ and $\mathbf{F} = \mathcal{F}$, that is f^0 is the function that minimizes the theoretical risk (2) in the class $\mathcal{F}$. Suppose that the loss function ϕ is "strictly convex" in the sense that

is described in Mendelson (2002); Bartlett et al. (2006) for the classification theory or Rejchel (2012) for the ranking problem. Omitting details that can be found in those papers we emphasize that many popular loss functions are strictly convex in that sense, for instance the exponential loss $\phi(t) = e^{-t}$, the truncated quadratic loss $\phi(t) = [\max(0, 1-t)]^2$ or the logistic loss $\phi(t) = \ln\left(1 + e^{-t}\right)$, but the hinge loss $\phi(t) = \max(0, 1-t)$ is not. In Rejchel (2012, Lemma 5) one proves that for such loss functions one obtains the condition (8) with the quadratic function G. Therefore, the convex conjugate H is also the quadratic function. Thus, the bound for the estimation risk seems to have the optimal rate. Obviously, the same result also holds for the pseudonorm $||\cdot||_c$. Finally, notice that the first two of above-mentioned convex loss functions do not satisfy the condition (7) from Assumption 1. They fulfill the Lipschitz condition on every compact subset of the real line, but their Lipschitz constants depend on these subsets. It could make our argumentation in the proof of Theorem 2 in Appendix A slightly more difficult. Of course, this inconvenience disappears, if the family $\mathcal{F}$ is uniformly bounded.

4 Conclusions

In the paper we study properties of estimators obtained by minimization of U-processes with the Lasso penalty. In the main theorem we obtained probabilistic inequalities concerning the quality of such estimators in prediction of the ordering between objects (via bounds for their excess risks) as well as model selection (via bounds for their l_1 distance from the oracle).

Besides, we should mention a possible improvement of results obtained in the paper, if different argumentation were used while proving the probabilistic inequality for a U-process in Appendix A. Namely, one could decompose a U-statistic into the sum of independent random variables and a degenerate U-statistic, and then bound these two terms separately. We do not proceed this way because of two reasons. First, following this argumentation we would meet some technical problems, for instance the contraction principle that plays a crucial role in the proof of Theorem 2 does not have an analogue for the Rademacher chaos of the order two that can be used to handle a degenerate U-process. The second reason is that this possible improvement would only replace the second moment of ψ_k in the definition of C^2 by its conditional version. It would have only a slight impact on the obtained results. For instance, this improvement would influence negligibly on rates of oracle inequalities in Theorem 1, especially comparing to the

impact of Assumption 2.

Finally, applying Theorem 1 to the weighted Lasso penalty need not significant changes in argumentation (van de Geer, 2008; Bühlmann and van de Geer, 2011). The weighted Lasso is a frequently used modification of its initial form, i.e. we minimize

$$Q_n(f_\theta) + \hat{\lambda}_n \sum_{k=1}^{m} w_k |\theta_k|,$$

where $w_1, \ldots, w_m$ are (possibly random) nonnegative weigths. Inserting weights into the penalty allows to consider parameters that are, in researcher's opinion, definitely significant in the model by setting their weights as zeros. Moreover, taking

$$w_k = \sqrt{\frac{1}{n(n-1)} \sum_{i \neq j} \psi_k^2(X_i, X_j)} \tag{18}$$

leads to the model with normalized base functions that is often used in practice. For instance, in the linear ranking model, i.e. $\psi_k(x, x') = x_k - x'_k$, the weight (18) is the empirical standard deviation of the k-th predictor X_k multiplied by $\sqrt{2}$. Therefore, using such weights we work with base functions that are measured on the same scale. Finally, our results cannot be applied to popular two-step algorithms, for instance the adaptive Lasso penalty (Zou, 2006). These procedures should be considered individually.

Appendix

The next two sections are devoted to the proof of Theorem 1.

A Stochastic arguments

In this section we tackle the stochastic part of the problem. It can be divided into two components. The first one relates to obtaining probabilistic inequalities for U-processes.

For an arbitrary number $M > 0$ define a family

$$\mathcal{F}_M = \{f_\theta \in \mathcal{F} : |\theta - \theta^*|_1 \leq M\}$$

and a U-process over a family $\mathcal{F}_M$

$$U(M) = \sup_{f_\theta \in \mathcal{F}_M} |Q_n(f_\theta) - Q_n(f_{\theta^*}) - Q(f_\theta) + Q(f_{\theta^*})|.$$

In this section we suppose that only Assumption 1 and 4 are satisfied. Henceforth we do not recall it.

Theorem 2. *With probability at least* $1 - \frac{1}{m^2}$

$$U(M) \leq 36\sqrt{3}\, ML\sqrt{\frac{\log m}{n}}\max(C, 6).$$

In the proof of Theorem 2 we use the following concentration inequality for U-processes which similar version can be found in Rejchel (2015, Theorem 2). We give its proof below to show the whole way of proving the main theorem that can be possibly improved as we have mentioned in Section 4.

Theorem 3. *Let* $\mathcal{G}$ *be a subset of a family of functions* $\{g : \mathcal{Z} \times \mathcal{Z} \to \mathbb{R}\}$ *that are uniformly bounded by a constant* $b > 0$ *and* $\sigma^2 = \sup_{g\in\mathcal{G}} \mathbf{Var} g(Z_1, Z_2)$. *Let an arbitrary* U*-process over a family* $\mathcal{G}$ *be given by*

$$U = \sup_{g\in\mathcal{G}} \left| \frac{1}{n(n-1)} \sum_{i\neq j} g(Z_i, Z_j) - \mathbf{E} g(Z_1, Z_2) \right|,$$

and for $N = \left\lfloor \frac{n}{2} \right\rfloor$ *we denote*

$$T = \sup_{g\in\mathcal{G}} \left| \frac{1}{N} \sum_{i=1}^{N} g(Z_i, Z_{N+i}) - \mathbf{E} g(Z_1, Z_2) \right|.$$

Then for every $t > 0$ *with probability at least* $1 - \exp(-nt^2)$

$$U \leq 2\mathbf{E}T + \sqrt{6}\sigma t + 10bt^2.$$

The expression T in Theorem 3 is the supremum of an average of independent random variables. Thus, this theorem reduces investigating properties of the U-process U to the significantly simpler empirical process T. Similar exponential inequalities for U-processes can be found in the literature, for instance Arcones (1995, Theorem 5). In fact, Theorem 3 is weaker than them, because it exploits the variance of functions in σ^2 instead of the conditional variance. However, the presented theorem is absolutely sufficient to obtain satisfactory bounds for $U(M)$ as we have discussed in Section 4. Moreover, its application to our problem is immediate.

Proof of Theorem 3. Fix $t, u, \lambda > 0$. From Markov's inequality we have

$$\begin{aligned} \mathbb{P}\left(U > \mathbf{E}T + \frac{u}{N}\right) &= \mathbb{P}\left[\exp\left(\lambda N U\right) > \exp\left(\lambda N \mathbf{E}T + \lambda u\right)\right] \\ &\leq \exp\left(-\lambda N \mathbf{E}T - \lambda u\right) \mathbf{E} \exp\left[\lambda N U\right]. \end{aligned} \tag{19}$$

We have mentioned in Section 1 that using Hoeffding's decomposition we can represent every U-statistic as an average of averages of independent random variables (Serfling, 1980), i.e.

$$\frac{1}{n(n-1)}\sum_{i\neq j} g(Z_i, Z_j) = \frac{1}{n!}\sum_{\pi}\frac{1}{N}\sum_{i=1}^{N} g\left(Z_{\pi(i)}, Z_{\pi(N+i)}\right), \tag{20}$$

where the first sum on the right-hand side of (20) is taken over all permutations π of a set $\{1,\dots,n\}$. Applying (Clémençon et al., 2008, Lemma A.1) we can estimate $\mathbf{E}\exp\left[\lambda NU\right]$ by $\mathbf{E}\exp\left(\lambda NT\right)$. Therefore, the right-hand side of (19) is bounded by

$$\exp(-\lambda u)\,\mathbf{E}\exp\left[\lambda N(T-\mathbf{E}T)\right]. \tag{21}$$

We have already mentioned that T is the supremum of an empirical process and concentration inequalities for it can be found in the literature (Talagrand, 1996; Massart, 2000; Boucheron et al., 2000; Bousquet, 2002). We use Bühlmann and van de Geer (2011, Theorem 14.1) and bound (21) by

$$\exp\left[-\lambda u + \left(\exp(2\lambda b) - 1 - 2\lambda b\right) v\right] \tag{22}$$

with $v = \left(\frac{N\mathbf{E}T}{b} + \frac{N\sigma^2}{4b^2}\right)$. Moreover, using argumentation from Boucheron et al. (2000, pages 281-282) we estimate (22) by

$$\exp\left[-\frac{3u^2}{4b(6bv+u)}\right].$$

Therefore, we have

$$\mathbb{P}\left(U > \mathbf{E}T + \frac{u}{N}\right) \leq \exp\left[-\frac{3u^2}{4b(6bv+u)}\right],$$

which easily implies that for every $x > 0$

$$\mathbb{P}\left(U \leq \mathbf{E}T + \sqrt{\frac{8bx}{N}\,\mathbf{E}T + \frac{2x\sigma^2}{N}} + \frac{4bx}{3N}\right) \geq 1 - \exp(-x).$$

To finish the proof of the theorem it is enough to take $x = nt^2$ and use two simple facts that $N \geq \frac{n}{3}$ and $2\sqrt{ab} \leq a + b$ for all nonnegative numbers a, b. □

Proof of Theorem 2. We apply Theorem 3 with $\mathcal{G} = \left\{\phi_{f_\theta} - \phi_{f_{\theta^*}} : f_\theta \in \mathcal{F}_M\right\}$. From the condition (7) and Assumption 4 we get $|g|_\infty \le ML\sqrt{\frac{n}{\log m}}$ and $\mathbf{Var} g \le M^2L^2C^2$ for every $g \in \mathcal{G}$. Thus, from Theorem 3 we obtain that for every $t > 0$ with probability at least $1 - \exp(-nt^2)$

$$U(M) \le 2\mathbf{E}T(M) + \sqrt{6}MLCt + 10ML\sqrt{\frac{n}{\log m}}\,t^2\,, \tag{23}$$

where

$$T(M) = \sup_{f_\theta \in \mathcal{F}_M} \left| \frac{1}{N}\sum_{i=1}^{N} \left[\phi_{f_\theta}(Z_i, Z_{N+i}) - \phi_{f_{\theta^*}}(Z_i, Z_{N+i})\right] - Q(f_\theta) + Q(f_{\theta^*}) \right|.$$

The expression $T(M)$ is again the supremum over an average of independent variables, so to bound its expectation we use standard methods from the empirical process theory. We start with Symmetrization Lemma (van der Vaart and Wellner, 1996, Lemma 2.3.1) that gives us

$$\mathbf{E}T(M) \le 2\,\mathbf{E} \sup_{f_\theta \in \mathcal{F}_M} \left| \frac{1}{N}\sum_{i=1}^{N} \varepsilon_i \left[\phi_{f_\theta}(Z_i, Z_{N+i}) - \phi_{f_{\theta^*}}(Z_i, Z_{N+i})\right] \right|, \tag{24}$$

where we have an additional sequence of iid random variables $\varepsilon_1, \ldots, \varepsilon_N$ (the Rademacher sequence). Variables ε_i's take values 1 or -1 with probability $\frac{1}{2}$ and are independent of the sample $Z_1, \ldots, Z_n$. Next, by the contraction principle in Ledoux and Talagrand (1991, Theorem 4.12) or Bühlmann and van de Geer (2011, Theorem 14.4) the right-hand side of (24) is bounded by

$$4L\,\mathbf{E} \sup_{f_\theta \in \mathcal{F}_M} \left| \frac{1}{N}\sum_{i=1}^{N} \varepsilon_i \left[f_\theta(X_i, X_{N+i}) - f_{\theta^*}(X_i, X_{N+i})\right] \right|,$$

that can be estimated by

$$4ML\,\mathbf{E} \max_{1 \le k \le m} \left| \frac{1}{N}\sum_{i=1}^{N} \varepsilon_i \psi_k(X_i, X_{N+i}) \right|. \tag{25}$$

Therefore, by Bernstein's inequality (van de Geer, 2008, Lemma A.1) we bound (25) by $8\sqrt{3}ML\sqrt{\frac{\log m}{n}}\left(C + \sqrt{3}\right)$. Finally, taking $t = \sqrt{2\frac{\log m}{n}}$ and using the inequality (23) we receive that with probability at least $1 - \frac{1}{m^2}$

$$U(M) \le 36\sqrt{3}\,ML\sqrt{\frac{\log m}{n}}\max(C, 6).$$

□

The second task of this section is proving that probability of the event $\left\{\frac{\lambda_n}{2} \leq \hat{\lambda}_n \leq 2\lambda_n\right\}$ is close to one. It follows from the next lemma that is similar to van de Geer and Tarigan (2006, Lemma 5.7). However, our estimator $\hat{C}^2$ is built on the basis of U-statistics that requires modified argumentation.

Lemma 1. *Each of the following two inequalities*

$$\frac{1}{2}\max(C, 6) > \max(\hat{C}, 6), \tag{26}$$

$$2\max(C, 6) < \max(\hat{C}, 6). \tag{27}$$

holds with probability at most $\frac{1}{m^2}$.

Proof. We start with the inequality (26). Similarly to the proof of van de Geer and Tarigan (2006, Lemma 5.7) we consider two cases. In the first one we assume that $C \leq 6$. Obviously, we have $\max(\hat{C}, 6) \geq \max(C, 6)$, that implies

$$\mathbb{P}\left(\frac{1}{2}\max(C, 6) > \max(\hat{C}, 6)\right) = 0.$$

In the second case we suppose that $C > 6$ and denote by ψ_{max} the base function with largest second moment, i.e. $\mathbf{E}\psi^2_{max}(X, X') = C^2$. Thus, probability of (26) is upper bounded by $\mathbb{P}\left(C > 2\hat{C}\right)$ that is not greater than

$$\mathbb{P}\left(\frac{1}{n(n-1)}\sum_{i\neq j}\psi^2_{max}(X_i, X_j) - C^2 < -\frac{3}{4}C^2\right). \tag{28}$$

From Bernstein's inequality for U-statistics (Serfling, 1980, Theorem A, Section 5.6.1) we can estimate (28) by

$$\exp\left(-\frac{\frac{9}{16}NC^4}{2C^2\frac{n}{\log m} + C^2\frac{n}{2\log m}}\right),$$

which is not greater than $\exp(-2.7\log m)$, because $C > 6$. It finishes the proof of (26).

Next, we work with the inequality (27). It is trivial, if $\hat{C} \leq 6$, therefore we suppose that $\hat{C} > 6$. Then $\max(\hat{C}, 6) = \hat{C}$ and we have the following

equality and two inequalities that can be easily argued

$$\mathbb{P}\left(2\max(C,6) < \max(\hat{C},6)\right) = \mathbb{P}\left(\hat{C}^2 - C^2 > 3\max(C^2,36)\right)$$
$$\leq \mathbb{P}\left(\max_k \frac{1}{n(n-1)}\sum_{i\neq j}\psi_k^2(X_i,X_j) - \mathbf{E}\psi_k^2 > 3\max(C^2,36)\right)$$
$$\leq \sum_{k=1}^{m}\mathbb{P}\left(\frac{1}{n(n-1)}\sum_{i\neq j}\psi_k^2(X_i,X_j) - \mathbf{E}\psi_k^2 > 3\max(C^2,36)\right).$$

Using again Bernstein's inequality for U-statistics we can bound the last expression by

$$m\exp\left(-\frac{9n\max(C^4,36^2)}{6C^2\frac{n}{\log m} + 6\max(C^2,36)\frac{n}{\log m}}\right),$$

that is, by simple calculations, not greater than $\frac{1}{m^2}$. □

B Deterministic arguments

In this section we present the proof of Theorem 1. It is based on Theorem 2 and Lemma 1 from Appendix A as well as standard argumentation concerning empirical risk minimization with the Lasso penalty (van de Geer and Tarigan, 2006; van de Geer, 2008; Bühlmann and van de Geer, 2011). However, we have to adjust the latter methods to the ranking problem.

Proof of Theorem 1. Let $M^* = \frac{16\varepsilon^*}{\lambda_n}$. From Theorem 2 and Lemma 1 we know that probability of the event

$$\{U(M^*) \leq \varepsilon_*\} \cap \left\{\frac{\lambda_n}{2} \leq \hat{\lambda}_n \leq 2\lambda_n\right\} \tag{29}$$

is not less than $1 - \frac{3}{m^2}$. Our argumentation in the rest of the proof is "deterministic". We start with defining $t = \frac{M^*}{M^* + |\hat{\theta}-\theta^*|_1}$ and $\tilde{\theta} = t\hat{\theta} + (1-t)\theta^*$. Obviuosly, we have $|\tilde{\theta}-\theta^*|_1 \leq M^*$, so $f_{\tilde{\theta}} \in \mathcal{F}_{M^*}$, because Θ is a convex subset of $\mathbb{R}^m$. Using convexity of the loss function ϕ we obtain the inequality

$$Q_n(f_{\tilde{\theta}}) \leq tQ_n(f_{\hat{\theta}}) + (1-t)Q_n(f_{\theta^*}),$$

which implies that

$$Q(f_{\tilde{\theta}}) - Q(f^0) \leq U(M^*) + t[Q_n(f_{\hat{\theta}}) - Q_n(f_{\theta^*})] + Q(f_{\theta^*}) - Q(f^0). \tag{30}$$

Combining the inequality (30) with facts that we consider the event (29) and $f_{\hat{\theta}}$ minimizes (5) we obtain the inequality

$$Q(f_{\tilde{\theta}}) - Q(f^0) + \hat{\lambda}_n|\tilde{\theta}|_1 \leq \varepsilon^* + \hat{\lambda}_n|\theta^*|_1 + Q(f_{\theta^*}) - Q(f^0).$$

Properties $|\theta|_1 = |\theta_{S_*}|_1 + |\theta_{S'_*}|_1$ and $\theta^*_{S'_*} = 0$ imply two inequalities

$$Q(f_{\tilde{\theta}}) - Q(f^0) + \hat{\lambda}_n|\tilde{\theta}_{S_*}|_1 + \hat{\lambda}_n|\tilde{\theta}_{S'_*}|_1 \leq \varepsilon^* + Q(f_{\theta^*}) - Q(f^0) + \hat{\lambda}_n|\theta^*_{S_*}|_1 \quad (31)$$

and

$$Q(f_{\tilde{\theta}}) - Q(f^0) + \hat{\lambda}_n|\tilde{\theta}_{S'_*}|_1 \leq 2\varepsilon^* + \hat{\lambda}_n|\tilde{\theta}_{S_*} - \theta^*_{S_*}|_1. \quad (32)$$

We will consider two cases. In the first one we assume that $\hat{\lambda}_n|\tilde{\theta}_{S_*} - \theta^*_{S_*}|_1 \leq \varepsilon^*$, which together with (32) and $\theta^*_{S'_*} = 0$ implies that

$$Q(f_{\tilde{\theta}}) - Q(f^0) + \hat{\lambda}_n|\tilde{\theta}_{S'_*} - \theta^*_{S'_*}|_1 \leq 3\varepsilon^*. \quad (33)$$

Adding $\hat{\lambda}_n|\tilde{\theta}_{S_*} - \theta^*_{S_*}|_1$ to both sides of (33) we get

$$Q(f_{\tilde{\theta}}) - Q(f^0) + \hat{\lambda}_n|\tilde{\theta} - \theta^*|_1 \leq 4\varepsilon^*, \quad (34)$$

so $\hat{\lambda}_n|\tilde{\theta} - \theta^*|_1 \leq 4\varepsilon^*$. Finally, on (29) we have $\lambda_n \leq 2\hat{\lambda}_n$, therefore we obtain that $|\tilde{\theta} - \theta^*|_1 \leq \frac{M^*}{2}$, so $|\hat{\theta} - \theta^*|_1 \leq M^*$.

Next, we consider the second case, i.e. we suppose that $\hat{\lambda}_n|\tilde{\theta}_{S_*} - \theta^*_{S_*}|_1 > \varepsilon^*$. Therefore, adding $\hat{\lambda}_n|\tilde{\theta}_{S_*} - \theta^*_{S_*}|_1$ to both sides of (31) and making simple calculations we obtain

$$Q(f_{\tilde{\theta}}) - Q(f^0) + \hat{\lambda}_n|\tilde{\theta} - \theta^*|_1 \leq \varepsilon^* + Q(f_{\theta^*}) - Q(f^0) + 2\hat{\lambda}_n|\tilde{\theta}_{S_*} - \theta^*_{S_*}|_1. \quad (35)$$

Besides, in this case we have from (32) that $\hat{\lambda}_n|\tilde{\theta}_{S'_*}|_1 \leq 3\hat{\lambda}_n|\tilde{\theta}_{S_*} - \theta^*_{S_*}|_1$, so $|\tilde{\theta}_{S'_*} - \theta^*_{S'_*}|_1 \leq 3|\tilde{\theta}_{S_*} - \theta^*_{S_*}|_1$. Therefore, we can use the compatibility condition for the set S_* (by Assumption 3) and bound the right-hand side of (35) by

$$\varepsilon^* + Q(f_{\theta^*}) - Q(f^0) + 2\hat{\lambda}_n \frac{||f_{\tilde{\theta}} - f_{\theta^*}||\sqrt{|S_*|}}{A_*}.$$

Using Assumption 2 as well as facts that $uv \leq G(u) + H(v)$ for every u, v and $\hat{\lambda}_n \leq 2\lambda_n$ we obtain the chain of inequalities

$$\begin{aligned}
2\hat{\lambda}_n \frac{||f_{\tilde{\theta}} - f_{\theta^*}||\sqrt{|S_*|}}{A_*} &\leq 4\delta \frac{\lambda_n\sqrt{|S_*|}}{\delta A_*}||f_{\tilde{\theta}} - f^0|| + 4\delta \frac{\lambda_n\sqrt{|S_*|}}{\delta A_*}||f_{\theta^*} - f^0|| \\
&\leq 8\delta H\left(\frac{\lambda_n\sqrt{|S_*|}}{\delta A_*}\right) + 4\delta G\left(||f_{\tilde{\theta}} - f^0||\right) + 4\delta G\left(||f_{\theta^*} - f^0||\right) \\
&\leq 8\delta H\left(\frac{\lambda_n\sqrt{|S_*|}}{\delta A_*}\right) + 4\delta\left[Q(f_{\tilde{\theta}}) - Q(f^0)\right] + 4\delta\left[Q(f_{\theta^*}) - Q(f^0)\right].
\end{aligned}$$

Summarizing, we can upper bound the left-hand side of the inequality (35) by

$$\varepsilon^* + (1+4\delta)\left[Q(f_{\theta^*}) - Q(f^0)\right] + 8\delta H\left(\frac{\lambda_n\sqrt{|S_*|}}{\delta A_*}\right) + 4\delta\left[Q(f_{\tilde{\theta}}) - Q(f^0)\right],$$

that gives us

$$(1-4\delta)\left[Q(f_{\tilde{\theta}}) - Q(f^0)\right] + \hat{\lambda}_n|\tilde{\theta} - \theta^*|_1 \leq 2\varepsilon^*. \tag{36}$$

The inequality (36) implies that $\hat{\lambda}_n|\tilde{\theta} - \theta^*|_1 \leq 2\varepsilon^*$, so $|\tilde{\theta} - \theta^*|_1 \leq \frac{M^*}{4}$.

Thus, in both cases we obtain that $|\hat{\theta} - \theta^*|_1 \leq M^*$ or to be more precise

$$\mathbb{P}\left(|\hat{\theta} - \theta^*|_1 \leq M^*\right) \geq 1 - \frac{3}{m^2}.$$

To finish the proof of Theorem 1 it is enough to repeat above argumentation with $\hat{\theta}$ instead of $\tilde{\theta}$ that leads to (34) or (36) with $\tilde{\theta}$ replaced by $\hat{\theta}$. □

Acknowledgements Research financed by Polish National Science Centre no. DEC-2014/12/S/ST1/00344.

References

Agarwal, S., Graepel, T., Herbrich, R., Har-Peled, S., and Roth, D. (2005). Generalization bounds for the area under the ROC curve. *J. Machine Learning Research*, 6:393–425.

Arcones, M. A. (1995). A Bernstein-type inequality for U-statistics and U-processes. *Statist. Probab. Letter*, 22:223–230.

Bartlett, P. L., Bousquet, O., and Mendelson, S. (2005). Local rademacher complexities. *Ann. Statist.*, 33:1497–1537.

Bartlett, P. L., Jordan, M. I., and McAuliffe, J. D. (2006). Convexity, classification and risk bounds. *Journal of the American Statistical Association*, 101:138–156.

Bickel, P. J., Ritov, Y., and Tsybakov, A. B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *Ann. Statist.*, 37:1705–1732.

Bobrowski, L. and Łukaszuk, T. (2012). Prognostic modeling with high dimensional and censored data. *Lecture Notes in Computer Science*, 7377:178–193.

Boucheron, S., Lugosi, G., and Massart, P. (2000). A sharp concentration inequality with applications. *Random Structures Algorithms*, 16:277–292.

Bousquet, O. (2002). A Bennett concentration inequality and its application to suprema of empirical processes. *C. R. Acad. Sci. Paris*, Ser. I 334:495–500.

Bühlmann, P. and van de Geer, S. (2011). *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer, New York.

Bunea, F., Tsybakov, A., and Wegkamp, M. H. (2007). Sparse density estimation with l_1 penalties. *COLT*, 4539:530–543.

Chen, H. and Wu, J. (2013). Regularized ranking with convex losses and l_1-penalty. *Abstract and Applied Analysis*, 2013.

Clémençon, S., Lugosi, G., and Vayatis, N. (2008). Ranking and empirical minimization of U-statistics. *Ann. Statist.*, 36:844–874.

Cortes, C. and Mohri, M. (2004). AUC optimization vs. error rate minimization. In *In Advances in Neural Information Processing Systems*, volume 16, pages 1436–1462.

Cossock, D. and Zhang, T. (2006). Subset ranking using regression. In *Proceedings of the 19th Annual Conference on Learning Theory*.

de la Peña, V. H. and Giné, E. (1999). *Decoupling: From Dependence to Independence*. Springer-Verlag, New York.

Freund, Y., Iyer, R., Schapire, R. E., and Singer, Y. (2004). An efficient boosting algorithm for combining preferences. *J. Machine Learning Research*, 4:933–969.

Greenshtein, E. (2006). Best subset selection, persistence in high dimensional statistical learning and optimization under l_1 constraint. *Ann. Statist*, 34:2367–2386.

Han, A. K. (1987). Non-parametric analysis of a generalized regression model. *Journal of Econometrics*, 35:303–316.

Huang, J. and Zhang, C.-H. (2012). Estimation and selection via absolute penalized convex minimization and its multistage adaptive applications. *Journal of Machine Learning Research*, 13:1839–1864.

Klein, R. W. and Spady, R. H. (1993). An efficient semiparametric estimator for binary response models. *Econometrica*, 61:387–421.

Lai, H., Pan, Y., Liu, C., Lin, L., and Wu, J. (2013). Sparse learning-to-rank via an efficient primal-dual algorithm. *IEEE Trans. Comput.*, 62:1221–1233.

Laporte, L., Flamary, R., Canu, S., Déjean, S., and Mothe, J. (2014). Non-convex regularizations for feature selection in ranking with sparse SVM. *IEEE Transactions on Neural Networks and Learning Systems*, 25:1118–1130.

Ledoux, M. and Talagrand, M. (1991). *Probability in Banach Spaces: Isoperimetry and Processes*. Springer, Berlin.

Mammen, E. and Tsybakov, A. B. (1999). Smooth discrimination analysis. *Ann. Statist.*, 27:1808–1829.

Massart, P. (2000). About the constants in Talagrand's inequality for empirical processes. *Ann. Probab.*, 29:863–884.

Meinshausen, N. and Buhlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *Ann. Statist*, 34:1436–1462.

Mendelson, S. (2002). Improving the sample complexity using global data. *IEEE Trans. Inform. Theory*, 48:1977–1991.

Negahban, S. N., Ravikumar, P., Wainwright, M. J., and Yu, B. (2012). A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science*, 27:538–557.

Rejchel, W. (2012). On ranking and generalization bounds. *Journal of Machine Learning Research*, 13:1373–1392.

Rejchel, W. (2015). Fast rates for ranking with large families. *Neurocomputing*, page http://dx.doi.org/10.1016/j.neucom.2015.05.013i.

Rudin, C. (2006). Ranking with a p-norm push. In *In Proceedings of the 19th Annual Conference on Learning Theory*.

Serfling, R. J. (1980). *Approximation Theorems of Mathematical Statistics*. Wiley, New York.

Sherman, R. P. (1993). The limiting distributions of the maximum rank correlation estimator. *Econometrica*, 61:123–137.

Talagrand, M. (1996). New concentration inequalities in product spaces. *Invent. Math.*, 126:503–563.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288.

van de Geer, S. (2008). High-dimensional generalized linear models and the lasso. *Annals of Statistics*, 36:614–645.

van de Geer, S. and Tarigan, B. (2006). Classifiers of support vector machine type with l_1 complexity regularization. *Bernoulli*, 12:1045–1076.

van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Verlag, New York.

Zhao, P. and Yu, B. (2006). On model selection consistency of lasso. *Journal of Machine Learning Research*, 7:2541–2563.

Zhou, S. (2009). Thresholding procedures for high dimensional variable selection and statistical estimation. *Advances in Neural Information Processing Systems*, 22:1436–1462.

Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101:1418–1429.